



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/700,227	11/03/2003	Steven R. Kleiman	112056-0143	3081
24267	7590	11/03/2006		
CESARI AND MCKENNA, LLP 88 BLACK FALCON AVENUE BOSTON, MA 02210			EXAMINER PATEL, HETUL B	
			ART UNIT 2186	PAPER NUMBER

DATE MAILED: 11/03/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

10/700,227

Applicant(s)

KLEIMAN ET AL.

Examiner

Hetul Patel

Art Unit

2186

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 03 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 23 October 2006.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-12 and 14-44 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-12 and 14-44 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
- ☐ Certified copies of the priority documents have been received.
 - ☐ Certified copies of the priority documents have been received in Application No. _____.
 - ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- ☐ Notice of References Cited (PTO-892)
- ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- ☐ Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date _____
- ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____
- ☐ Notice of Informal Patent Application
- ☐ Other: _____

DETAILED ACTION

Continued Examination Under 37 CFR 1.114

1. A request for continued examination under 37 CFR 1.114, including the fee set forth in 37 CFR 1.17(e), was filed in this application after final rejection. Since this application is eligible for continued examination under 37 CFR 1.114, and the fee set forth in 37 CFR 1.17(e) has been timely paid, the finality of the previous Office action has been withdrawn pursuant to 37 CFR 1.114. Applicant's submission filed on 10/23/2006 has been entered and carefully considered.
2. Applicant's arguments filed on 10/23/2006 have been fully considered but they are not persuasive.

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

3. Claims 1-5, 22-25 and 28-33 are rejected under 35 U.S.C. 102(b) as being anticipated by "A case for Redundant Arrays of Inexpensive Disks (RAID)" by Patterson et al. hereinafter, Patterson.

As per claims 1-3, Patterson teaches a system adapted to distribute redundant information across disks of an array, the system comprising: a storage operating system

configured to invoke storage operations executed by a storage system, the storage operating system further configured to manage storage of information, including the redundant information and data, on blocks of the disks in response to disk access operations, the storage operating system including a storage module adapted to compute the redundant information in response to placement of the data in stripes across the disks, the storage operating system maintaining at least one unallocated block (i.e. one block for check/parity data) per stripe for use by the storage module to store the computed redundant information, wherein the at least one unallocated block used to store the redundant information is located in any disk (i.e. since the check/parity data are written in every disk in round robin/circular pattern) and wherein the location of the at least one unallocated block use to store the redundant information is dynamically allocated by the storage module before each write request for each stripe (i.e. the check data is getting written to different disk each time in round robin pattern as shown in Fig. 4b); and the disk array controller and the RAID system are configured to compute the redundant information and reconstruct blocks lost due to failure of one or more of the disks (e.g. see "Fifth Level RAID" section 11 on page 114 and Fig. 4).

As per claims 22, 31 and 33, see arguments with respect to rejection of claim 1. Claims 22, 31 and 33 are also rejected based on the same rationale as the rejection of claim 1.

Regarding claims 24 and 29-30, Patterson, in his paper, describes distributing redundant info across disks of an array, dividing each disk into blocks (which are organized into stripes and each stripe contains one block from each disk), selecting any

block in the stripe not used to contain data to contain redundant information, wherein the block used is located in any disk, where in the location of the block can change over time, computing redundant information based on contents of all blocks in the stripe, assigning the block to contain redundant information when each stripe is written, and computing redundant information using algebraic and algorithmic calculations in response to placement of data on the array.

The examiner notes that the above claims describe a level 5 RAID system. Patterson's paper teaches about all of the limitations given in the claims above and, although his paper does not expressly state so, many of the other claimed limitations.

Regarding claims 4, 5, 23, and 32, a storage operating system further configured to implement a high level module that maintains information about locations of data on the disks, the high level module being a file system or database adapted to control the layout of data on the disks, and determining which block in the stripe contains redundant information each time there is a write request to the stripe is found in the Microsoft Computer Dictionary and is used with evidentiary support. The dictionary states:

"A table or list maintained by some operating systems to manage disk space used for file storage. Files on a disk are stored, as space allows, in fixed-size groups of bytes (characters) rather than from beginning to end as contiguous strings of text or numbers. A single file can thus be scattered in pieces over many separate storage areas. A file allocation table maps available disk storage space so that it can mark

flawed segments that should not be used and can find and link pieces of a file. In MS-DOS, the file allocation table is commonly known as the FAT.”

Regarding claims 25 and 28, the storage module being one of an array controller and RAID system is taught by Patterson.

The step of determining is performed by a high level module, the module being a file system, of a storage system and wherein the steps of computing and assigning are performed by a storage module of the storage system is understood by the Microsoft Computer Dictionary excerpt above and the ability for the Microsoft Windows 2000 operating system to do all the functions of a level 5 RAID array.

4. Claims 36, 39, 42, and 44 rejected under 35 U.S.C. 102(e) as being clearly anticipated by Spiegel et al. (Pat No 6571326).

As per claims 36, 39, 42, and 44, Spiegel et al. teaches determining which blocks are unallocated in a strip across the disks; reserving unallocated blocks for storing the redundant information in one or more reserved unallocated blocks; arranging data in the stripe for the data to be stored in one or more allocated blocks across the disks of the array; assigning the redundant information to the one or more reserved unallocated blocks; and writing the data in the allocated blocks and the redundant information in the one or more reserved unallocated blocks as the stripe across the disks of the array (column 4 line 44 to column 6 line 20).

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

5. Claims 6-8, 11-12 and 26 rejected under 35 U.S.C. 103(a) as being unpatentable over Patterson in view of Baylor et al. (Pat No 5862158).

Regarding claims 6-8, Patterson and evidentiary support describes all of the limitations found in claims 1-5. The file system or database configured to determine block locations of data and redundant info on disks is found in the Microsoft Computer Dictionary and is used with evidentiary support. The dictionary states:

"A table or list maintained by some operating systems to manage disk space used for file storage. Files on a disk are stored, as space allows, in fixed-size groups of bytes (characters) rather than from beginning to end as contiguous strings of text or numbers. A single file can thus be scattered in pieces over many separate storage areas. A file allocation table maps available disk storage space so that it can mark flawed segments that should not be used and can find and link pieces of a file. In MS-DOS, the file allocation table is commonly known as the FAT."

Patterson and evidentiary support do not teach the storage operating system integrates the file system or database with the RAID system, the file system or database configured to determine block locations of the data on disks, or the RAID system configured to determine block locations of redundant information on the disks.

Baylor et al. discloses the storage operating system integrates the file system or database with the RAID system (column 10, lines 35-46), the file system or database configured to determine block locations of the data on disks (this capability found within Windows 2000 and used with evidentiary support), and the RAID system configured to determine block locations of redundant information on the disks (column 7, lines 12-22, table 10, and table 20).

Patterson and Baylor et al. are analogous art because they are from a similar problem solving area, storage arrays and redundancy. At the time of the invention it would have been obvious to a person of ordinary skill in the art to integrate the storage system and have it know the location of all blocks. The suggestion for doing so would have been access time and a speed increase. Therefore, it would have been obvious to combine Baylor et al. and Patterson for the benefit of speed to obtain the invention as specified in claims 6-8.

The examiner notes that the FAT keeps locations of all data and that claim 8, as written, is defined broadly enough for the FAT to describe the limitations found in claim 8. The examiner also notes that because the functionality of a RAID in Windows 2000, the file system and storage operating system are highly integrated and with Windows 2000 using a file system (FAT being the prime example), it would also be configured to know the block locations of data on the disks.

Regarding claim 11, the block allocation map structures used by the file system to determine block locations of data and redundant information on disks is found in the

Microsoft Computer Dictionary and is used with evidentiary support. The dictionary states:

"A table or list maintained by some operating systems to manage disk space used for file storage. Files on a disk are stored, as space allows, in fixed-size groups of bytes (characters) rather than from beginning to end as contiguous strings of text or numbers. A single file can thus be scattered in pieces over many separate storage areas. A file allocation table maps available disk storage space so that it can mark flawed segments that should not be used and can find and link pieces of a file. In MS-DOS, the file allocation table is commonly known as the FAT."

The examiner notes that the FAT keeps locations of all data and that claim 11, as written, is defined broadly enough for the FAT to describe the limitations found in claim 11.

Regarding claim 12, Patterson describes the redundant information being parity throughout his original paper.

Regarding claim 26, Patterson and evidentiary support describe all of the limitations found in claims 22-25. Patterson also discloses maintaining at least one unallocated block per stripe for use by the storage module.

Patterson does not expressly disclose providing indication from a high level module to the storage module of an unallocated block to contain parity.

Stallmo et al. does disclose providing indication from a high level module to the storage module of an unallocated block to contain parity (column 7, lines 57-65).

Patterson and Stallmo et al. are analogous art because they are from a similar problem solving area, storage arrays and redundancy. At the time of the invention it would have been obvious to a person of ordinary skill in the art to provide indication of which block will be parity. The suggestion for doing so would have been faster recovery. Therefore, it would have been obvious to combine Stallmo et al. and Patterson for the benefit of faster recovery to obtain the invention as specified in claim 26.

6. Claims 14-15, 18-20, 27 and 34-35 rejected under 35 U.S.C. 103(a) as being unpatentable over Patterson and Baylor et al., and further in view of Stallmo et al. (Pat No 5657468).

Regarding claim 14, Patterson describes in his paper all the limitations in claim 1 and also describes selecting at least one unallocated block to store redundant information.

Patterson does not teach expressly the storage module computing redundant information using the redundant storage algorithm.

Stallmo et al. does disclose the storage module computing redundant information using the redundant storage algorithm (column 3, lines 5-35).

Patterson and Stallmo et al. are analogous art because they are from a similar problem solving area, storage arrays and redundancy. At the time of the invention it would have been obvious to a person of ordinary skill in the art to compute the redundant info using the storage algorithm. The suggestion for doing so would have

been efficiency. Therefore, it would have been obvious to combine Stallmo et al. and Patterson for the benefit of efficiency to obtain the invention as specified in claim 14.

Regarding claim 15, Patterson and Stallmo et al. describe all the limitations found in claims 1 and 14, but do not expressly describe a selection of at least one unallocated block is independent of the redundant storage algorithm.

Baylor et al. does disclose a selection of at least one unallocated block is independent of the redundant storage algorithm (column 7, lines 12-22, table 10, and table 20).

Patterson, Stallmo et al. and Baylor et al. are analogous art because they are from a similar problem solving area, storage arrays and redundancy. At the time of the invention it would have been obvious to a person of ordinary skill in the art to have the selection of the location of the parity block not be by the redundant storage algorithm. The suggestion for doing so would have been because the file system or operating system would be better suited for that function. Therefore, it would have been obvious to combine Baylor et al. Stallmo et al. and Patterson for the benefit of efficiency to obtain the invention as specified in claim 15.

Furthermore, Stallmo et al. teaches that the redundant storage algorithm is an asymmetric algorithm and wherein the redundant information is parity (column 3, lines 5-35) as per claims 34-35.

Regarding claims 18-20, Patterson and Stallmo et al. describe all the limitations found in claims 1 and 14. However, neither describe at least one unallocated block used to store redundant information comprises two or more unallocated blocks used to store

redundant information, a selection of unallocated blocks to store redundant information is independent of the redundant storage algorithm used to compute redundant information, or the redundant storage algorithm depends on positions of blocks in the array.

Baylor et al. does describe at least one unallocated block used to store redundant information comprises two or more unallocated blocks used to store redundant information (figures 2 and 3), a selection of unallocated blocks to store redundant information is independent of the redundant storage algorithm used to compute redundant information, and the redundant storage algorithm depends on positions of blocks in the array (column 7, lines 12-22, table 10, and table 20).

Patterson, Stallmo et al. and Baylor et al. are analogous art because they are from a similar problem solving area, storage arrays and redundancy. At the time of the invention it would have been obvious to a person of ordinary skill in the art to have two parity blocks placed independently of the redundant storage algorithm. The suggestion for doing so would have been to make sure that recovery is certain. Therefore, it would have been obvious to combine Baylor et al., Stallmo et al. and Patterson for the benefit of certain recovery to obtain the invention as specified in claims 18-20.

Regarding claim 27, Patterson, Stallmo et al., and evidentiary support describe all of the limitations found in claims 22-26. These prior arts do not expressly disclose, reconstructing, using the storage module, a block that is lost due to disk failure.

Baylor et al. does disclose, reconstructing, using the storage module, a block that is lost due to disk failure (column 10, lines 19-30).

Patterson, Stallmo et al. and Baylor et al. are analogous art because they are from a similar problem solving area, storage arrays and redundancy. At the time of the invention it would have been obvious to a person of ordinary skill in the art to reconstruct lost data. The suggestion for doing so would have been to not lose important information. Therefore, it would have been obvious to combine Baylor et al., Stallmo et al., and Patterson for the benefit of data protection to obtain the invention as specified in claim 27.

The examiner notes that Patterson does talk about recovery, but that the prior art of Baylor et al. goes into much greater detail of data recover through parity stripes.

7. Claims 9 and 10 rejected under 35 U.S.C. 103(a) as being unpatentable over Patterson and Baylor et al., and further in view of Ulrich et al. (Pub No 20020124137).

Patterson and Baylor et al. describe all the limitations of claims 1, 3-6 and 8, but do not expressly disclose the file system or database rendering balancing decisions to determine block locations of data and redundant information on disks and balancing decisions comprising one of different sizes and speeds of disks and whether the disk is more utilized than others.

Ulrich et al. discloses the file system or database rendering balancing decisions to determine block locations of data and redundant information on disks and balancing decisions comprising one of different sizes and speeds of disks and whether the disk is more utilized than others (Paragraph 30).

Patterson, Baylor et al., and Ulrich et al. are analogous art because they are from a similar problem solving area, storage arrays and redundancy. At the time of the invention it would have been obvious to a person of ordinary skill in the art to do balancing decisions on the volume. The suggestion for doing so would have been speed and equal wear. Therefore, it would have been obvious to combine Ulrich et al, Baylor et al., and Patterson for the benefit of speed and equal wear to obtain the invention as specified in claims 9 and 10.

8. Claim 16, 17, and 21 are rejected under 35 U.S.C. 103(a) as being unpatentable over Patterson, Baylor et al., and Stallmo et al., and further in view of Wiencko et al. (Pat No 6557123).

Patterson, Baylor et al., and Stallmo et al. describe all the limitations of claims 1, 14, 15, and 18-20. Patterson teaches the redundant information being parity. Patterson Baylor et al., and Stallmo et al. do not expressly disclose that the redundant storage algorithm is a symmetric or asymmetric algorithm.

Wiencko et al. does disclose redundant information being parity. Patterson Baylor et al., and Stallmo et al. do not expressly disclose that the redundant storage algorithm is a symmetric or asymmetric algorithm (column 9 line 57 to column 10 line 2).

Patterson, Baylor et al., Stallmo et al., and Wiencko et al. are analogous art because they are from a similar problem solving area, storage arrays and redundancy. At the time of the invention it would have been obvious to a person of ordinary skill in the art to use a symmetric algorithm for the redundant storage. The suggestion for doing

so would have been for reversibility. Therefore, it would have been obvious to combine Wiencko et al., Baylor et al., Stallmo et al., and Patterson for the benefit of reversibility to obtain the invention as specified in claims 16, 17, and 21.

9. Claims 37-38, 40-41 and 43 are rejected under 35 U.S.C. 103(a) as being unpatentable over Spiegel et al. in view of Paterson.

Spiegel et al. describes storing a second stripe across the array by determining one or more unallocated blocks across the array (column 4 line 44 to column 6 line 20).

Spiegel et al. does not describe storing parity information as redundant information in a reserved unallocated block, adding another disk to the array, writing the data to the allocated blocks and the redundant information to the unallocated blocks of the second stripe. Patterson describes all this functionality.

Spiegel et al. and Patterson are analogous art because they are from the same field of endeavor, data allocation in a multi drive system. At the time of the invention it would have been obvious to a person of ordinary skill in the art to allow for more drives and multiple stripes. The suggestion for doing so would have been speed increases. Therefore, it would have been obvious to combine Patterson and Spiegel et al. for the benefit of speed to obtain the invention as specified in claims 37, 38, 40, 41, and 43.


Conclusion

10. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Hetul Patel whose telephone number is 571-272-4184. The examiner can normally be reached on M-F 8-4:30:

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matt Kim can be reached on 571-272-4182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

HBP
HBP


PIERRE BATAILLE
PRIMARY EXAMINER
10/30/06